Mechanism Design

November 29, 2022

One problem that has been mentioned throughout the class is that certain information, such as values purchasers have for items or costs producers have to supply items, are not known. This lack of knowledge makes it difficult to determine how well our models predict behavior. The goal of mechanism design is to elicit that information from individuals. The problem is that individuals may not truthfully reveal such information, particularly if the incentives are structured such that truthful revelation is not optimal. One may think it morally wrong to not truthfully reveal one's value for an item when asked, but when people purchase items from a store they are not necessarily revealing their maximum willingness to pay for a certain good, just that they are willing to pay a certain price (the posted price) for the good.

At the beginning of the course we discussed that the fundamental problem in economics is how to allocate scarce resources. We discussed a few different methods for rationing goods – lottery, first come first serve, merit, randomly, etc. – but that our focus would be on the price system. There are many different ways that we can use prices to allocate goods. We have posted-offer markets, negotiations, pit markets, auctions, etc. The goal now is to focus on a particular type of allocation system, which we will call a "mechanism."

Consider a game of incomplete information, where one party in a transaction does not know some information (perhaps the willingness to pay) of the other party has. Take a moment to consider what this means – there have to be multiple "types" of individuals in this model; if everyone was the same then there would be no mystery about anyone's preferences. At least one party does not know the "type" of the other party.[1] That means that one party could misrepresent their type to the other (state that they have a low value for an item when in fact they have a high value for an item).

Mechanism design can generally be thought of in three steps:

1. A seller designs a mechanism/contract/incentive scheme in which buyers send "messages" and allocations are made based upon the messages sent. An "allocation" here refers not only to the actual good(s) being transferred, but also any transfer payments that are made between buyer(s) and seller. Our focus will be on messages sent at the same time.

2. Buyers choose whether or not they want to participate in the game, or, alternatively one could think about this step as the buyers accepting or rejecting the seller's proposed mechanism. If buyers choose not to participate in the game they would typically have some reservation level of utility. This step requires consideration of the buyer's participation constraint (PC) when the buyer determines whether or not to participate.[2] There are some instances in which the "seller" can force the "buyer" to participate, thus removing the need for a participation constraint. Consider the government imposing a tax scheme on individuals or corporations. These economic agents must participate, even if "participating" means "not filing taxes properly and suffering the consequences." That scenario is different than eBay, which has no authority to force individuals to bid in its auctions or punish them for failing to bid in its auctions.

3. The final step is that buyers who choose to participate in the game submit their messages, and then an allocation is made based upon the submitted messages and the allocation rules of the mechanism.

Another key constraint that the seller must consider when designing the mechanism is the buyer's incentive compatibility (IC) constraint. Sellers want buyers to behave in a certain manner, and a properly

[1] Everyone knows their own type, so a particular individual has no uncertainty about their own type.

[2] Some texts refer to this constraint as the individual rationality (IR) constraint.

designed mechanism can elicit certain types of responses from buyers. Alternatively, in a principal-agent framework, the principal would like to structure a contract to align the incentives of the agent with those of the principal.

1 Background Results

There are two fairly general results that apply to broad classes of mechanisms. The first result is the Revelation Principle and the second result is the Revenue Equivalence Theorem. These two results can be thought of as corresponding to efficiency goals and revenue goals.

1.1 Revelation Principle

Suppose that a seller is concerned with efficiency, meaning that the highest valued user receives the good. While businesses are typically only concerned with which individual will pay the most for a good, governments may be concerned that the person who values the good the most receives the good. Suppose that a seller wishes to sell an object using some mechanism – the details of the mechanism are left unspecified as long as the following conditions are met:

1. The buyers simultaneously make claims about their types. Buyer i can claim to be any type from the possible set of types.
2. Given the buyers' claims, buyer i pays an amount that is a function of all the reported types and receives the good with some probability based upon the reported types

Games that satisfy these criteria are known as direct mechanisms, because the only action is to submit a claim about a type. Note that there is no claim as to whether that claim has to be truthful. An instructor could ask students how much effort they put into a project and then assign a prize based on reported effort. That would likely not be a very good mechanism at allocating the prize efficiently (if efficient means to the student with the highest actual effort), but it is a direct mechanism.

Some mechanisms have very complicated equilibrium strategies and it may be difficult to determine the equilibrium when individuals are not using truth-telling strategies. However, the revelation principle states that if an equilibrium exists in a game, then we can find a direct, truth-telling mechanism that has the same general properties. That is useful because if there is a mechanism in which individuals reveal their true preferences, then we can determine if the mechanism is efficient (meaning the individual(s) with the highest value(s) receive the good(s)). These are some formal and informal statements of the revelation principle for various types of mechanisms.[3] Note that a *Bayes-Nash equilibrium* has the same basic feature as a Nash equilibrium (set of strategies such that no player can unilaterally deviate from the set of strategies and receive a strictly higher payoff), only now each "type" must be playing a best response.

The following are some more formal statements of the revelation principle.

Proposition 1 *(Mas-Colell, Whinston, and Green pg. 493) Denote the set of possible states by Θ. In searching for an optimal contract, the owner can without loss restrict himself to contracts of the following form:*

1. *After the state θ is realized, the manager is required to announce which state has occurred.*
2. *The contract specifies an outcome $\left[w\left(\widehat{\theta}\right), e\left(\widehat{\theta}\right)\right]$ for each possible announcement $\widehat{\theta} \in \Theta$. (Note that w and e are just wages and efforts.)*
3. *in every state $\theta \in \Theta$, the manager finds it optimal to report the state truthfully.*

Proposition 2 *(Mas-Colell, Whinston, and Green pg. 884) Suppose that there exists a mechanism $\Gamma = (S_1, ..., S_I, g(\cdot))$ that implements social choice function $f(\cdot)$ in Bayesian Nash equilibrium. Then $f(\cdot)$ is truthfully implementable in Bayesian Nash equilibrium.*

[3] See Myerson (1979) for more detail.

Proposition 3 *(Fudenberg and Tirole, pg. 255) The principal can content herself with "direct" mechanisms, in which the message spaces are the type spaces, all agents accept the mechanism in step 2 regardless of their types, and the agents simultaneously and truthfully announce their types in step 3.*

Proposition 4 *(Gibbons, pg. 165) Any Bayesian Nash equilibrium of any Bayesian game can be represented by an incentive-compatible direct mechanism.*

Proposition 5 *(Wolfstetter, pg. 214) For any equilibrium of any auction game, there exists an equivalent incentive-compatible direct auction that leads to the same probabilities of winning and expected payments.*

1.2 Revenue Equivalence

All sellers, even the government, which might have additional concerns such as efficiency and equity, are concerned about revenue generation. The specific focus here will be on revenue generated from auction mechanisms but the results can be generalized to other mechanisms. Assume that all bidders have a value that is drawn from the same distribution of possible values. This distribution might be a uniform distribution over a particular range (say $[0, 100]$) or it could be from a normal distribution with some particular mean and variance.[4] The key is that all bidder values are drawn from the same distribution – that assumption is called *symmetric*. Second, the draws from the distribution are statistically *independent* so one draw does not influence the other. Third, the draws from the distribution are *private*, so each individual only knows their own private draws. Finally, we assume that the individuals all have *risk neutral* preferences. These assumptions constitute a symmetric, independent, private value environment with risk-neutral agents, which is the basis for our result. Suppose there is one object for sale.[5]

If the above conditions are met and the following two conditions:

1. The object always goes to the buyer with the highest value
2. any buyer with the lowest possible value draw expects 0 surplus

then any mechanism that satisfies the above assumptions yields the same expected revenue and results in a buyer with value v_i making the same expected payment across all mechanisms.

This result relies upon the previous result (revelation principle). The general idea is that there are direct mechanisms that are not truth-telling mechanisms, but they have the same Bayes-Nash equilibrium as a truth-telling mechanism by the revelation principle. We can then compare the expected revenue from the truth-telling mechanisms because the expected revenue relies on the underlying probability distribution of values.

It is important to consider all the assumptions that are used in deriving that result. Buyers are risk-neutral. Values are drawn from the same probability distribution. Value draws are statistically independent. While these assumptions are restrictive, they establish a benchmark for comparing more realistic settings. What happens if buyers are not risk-neutral? What happens if the distribution is not symmetric? What happens if value draws are not independent? What happens if value draws are not private? We will discuss these later because the predictions about revenue change depending upon which assumptions do not hold.

A more formal statement about the expected revenue from auctions follows:

Proposition 6 *(Mas-Colell, Whinston, and Green) 23.D.3 (Revenue Equivalence Theorem) Consider an auction setting with I risk-neutral buyers, in which buyer i's valuation is drawn from an interval* $\left[\underline{\theta}_i, \overline{\theta}_i\right]$ *with* $\underline{\theta}_i \neq \overline{\theta}_i$ *and a strictly positive density* $\phi_i(\cdot) > 0$*, and in which buyer's types are statistically independent. Suppose that a given pair of Bayesian Nash equilibria of two different auction procedures are such that for every buyer* i *: (i) For each possible realization of* $(\theta_i, ..., \theta_I)$*, Buyer i has an identical probability of getting the good in the two auctions; and (ii) Buyer i has the same expected utility level in the two auctions when his valuation for the object is at its lowest possible level. Then these equilibria of the two auctions generate the same expected revenue for the seller.*

[4] There are more formal requirements such that the distribution is strictly increasing and atomless. Strictly increasing simply means that the probability of drawing a value less than a number X must be less than the probability of drawing a value less than the number Y if $Y > X$. Atomless is similar to continuous, and really means that the probability of drawing a particular value is 0 (because we are focusing on drawing one particular value out of an infinite number of possible values).

[5] This result can be extended to the multiple object case as long as no buyer wants more than 1 of k identical, indivisible objects for sale

2 Mechanism Design Examples

The principles of mechanism design can be applied in a variety of settings. It could be one individual setting up an incentive system for another individual, such as in a principal-agent problem. It could be a government attempting to sell rights to producers of goods. It could be trying to match individuals with goods. It could be determining who the governing individuals are. We will discuss a principal-agent problem as a motivating example, then auctions, matching, and voting.

2.1 Principal-Agent Problem

The principal-agent problem occurs when the owner of a business (the principal) hires a manager (the agent) to run the business. For corporations this situation is usually the shareholders hiring a CEO, but could represent any manager and employee relationship, including those in government agencies or policy arenas. The problem is that some part of the agent behavior (typically labeled effort) is unobservable to the principal and that behavior tends to be correlated with productivity. However, productivity could also be influenced by factors that are not under the agent's control, so it is possible that the agent exerts high effort but, due to bad luck, productivity is low. The goal of the principal is to design a contract (or mechanism) to incentivize the agent to deliver high effort. The main goal of this example is to illustrate the participation and incentive compatibility constraints discussed earlier.

2.1.1 Game Tree (Extensive Form)

In our discussion of oligopoly we discussed some basic simultaneous games in which both firms chose a strategy at the same time. There are also sequential games (like the chess example) in which players take turns making decisions. Sequential games can be represented in a game tree (also called the extensive form). Game trees consist of the following pieces. There is an *initial node* to the game tree, which is the starting point of the game (think about the setup of the board when a game of Chess is begun – this is the initial node). From that initial node there are actions that the first mover can take. These actions are represented as *branches* to the game tree. At the end of each branch is a node. If the first mover makes a move and the game ends after that move, then we say that the game has reached a terminal node. A *terminal node* is a node at which no more actions can be taken. If the first mover makes a move and the second mover then gets to choose an action, we call this a *decision node*. The second mover's actions are then represented by branches extending from the decision node. The game tree extends until all the nodes are terminal nodes. At the terminal nodes, the payoffs to the players are listed. It is the convention to list the payoffs in the order that the players moved. One other important aspect of the game tree is the *information set*. For the initial principal-agent game we will initially consider, all decision nodes will also be information sets. However, it is possible that a game is being played and a player is uncertain as to which of a few decision nodes the player is at. In this case, the collection of decision nodes is that player's information set.

Solving Games in Game Tree Form One method of solving sequential games is to use the process of backward induction. This process simply means to start at the end of the game and work towards the beginning. So we would begin with the player who makes the last decision, determine what that player would do at the last decision node, and then work backwards with the other players knowing what future players would do at their decision nodes. Essentially we are "lopping off branches of the tree." This process removes noncredible threats by players earlier in the game and focuses on what a specific player would do once their decision node is reached, regardless of the path that led to that decision node being reached.[6]

2.1.2 Principal-Agent Observable Effort Case

We begin the discussion of the principal-agent problem by assuming that the agent's effort is observable. The game begins at an initial node with the principal deciding whether or not to make a contract offer to the agent. If the principal does not make a contract offer then the principal receives 0 and the agent

[6] The backward induction process will identify a set of strategies that is a subgame perfect Nash equilibria. Any subgame perfect Nash equilibria are Nash equilibria, but not all Nash equilibria are subgame perfect. This distinction is not important for the games we will discuss but is important in other games.

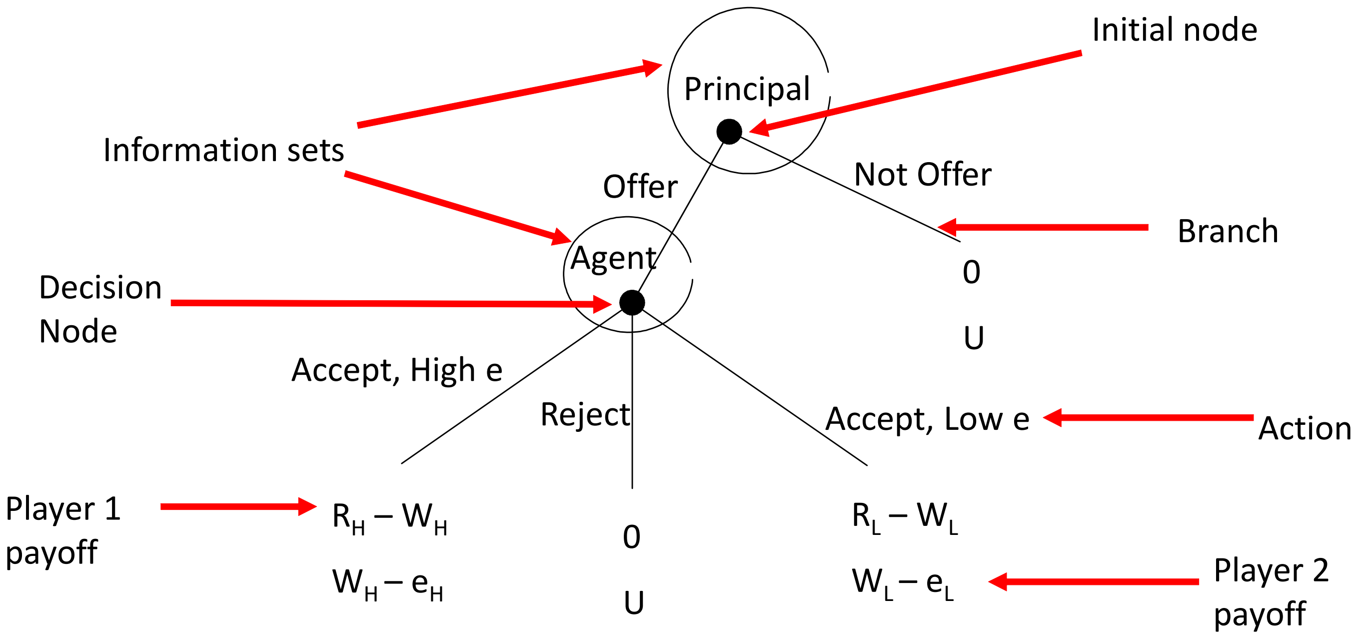

Figure 1: The game tree for the principal-agent problem with observable agent effort, with game components labeled.

receives some reservation utility U. If the principal makes a contract offer then it is the agent's move. The principal's goal is to determine what contract offer (which will be a set of wages in this game) will maximize profit. The contract specifies that the agent will receive W_H if high effort is exerted and W_L if low effort. The agent has three choices: accept the contract and exert high effort, accept the contract and exert low effort, or reject the contract. If the agent rejects the contract then the principal receives 0 and the agent receives U. If the agent accepts the contract and exerts low effort then the agent receives $W_L - e_L$, where e_L is the effort cost of exerting low effort and the principal receives $R_L - W_L$, where R_L is the revenue the principal receives if the agent exerts low effort. If the agent accepts the contract and exerts high effort then the agent receives $W_H - e_H$, where e_H is the effort cost of exerting high effort and the principal receives $R_H - W_H$, where R_H is the revenue the principal receives if the agent exerts high effort. Assume that e_H, e_L, R_H, and R_L are known; that the cost of exerting high effort is greater than that of exerting low effort, so $e_H > e_L$; and that the revenue from high effort is greater than the revenue from low effort, so $R_H > R_L$. Again, note that the agent's effort level is perfectly observable in this game. Figure 1 provides the game tree with the components identified; Figure 2 provides a "clean" version of the game tree.

The question to then ask is what restrictions are necessary on the parameters and the wage offer to have a Nash equilibrium of the game such that the principal offers the contract and the agent accepts and exerts high effort? From the principal's point of view, we need $R_H - W_H > 0$ because if it is not then the principal could be better off choosing to not offer the contract. It would also be helpful if $R_H - W_H > R_L - W_L$ because then the principal would prefer the agent to exert high effort. If the principal prefers that the agent exert low effort then it is fairly easy to ensure this by simply setting $W_H = W_L$ so that the agent receives the same payment regardless of which effort level is chosen. From the agent's point of view, two things need to happen. One is that $W_H - e_H > W_L - e_L$, so that the agent finds it more profitable to exert high effort rather than low effort. It also needs to be the case that $W_H - e_H > U$ so that the agent chooses to accept the contract rather than reject the contract. It is possible that $W_H - e_H > W_L - e_L$ but that exerting effort for this agent is too costly relative to the agent's opportunity cost (U) so the agent would simply choose to reject the contract. As we have already discussed with mechanisms, these two conditions are just the incentive compatibility constraint and the participation constraint.

Incentive compatibility constraint: The principal must structure the contract such that it gives

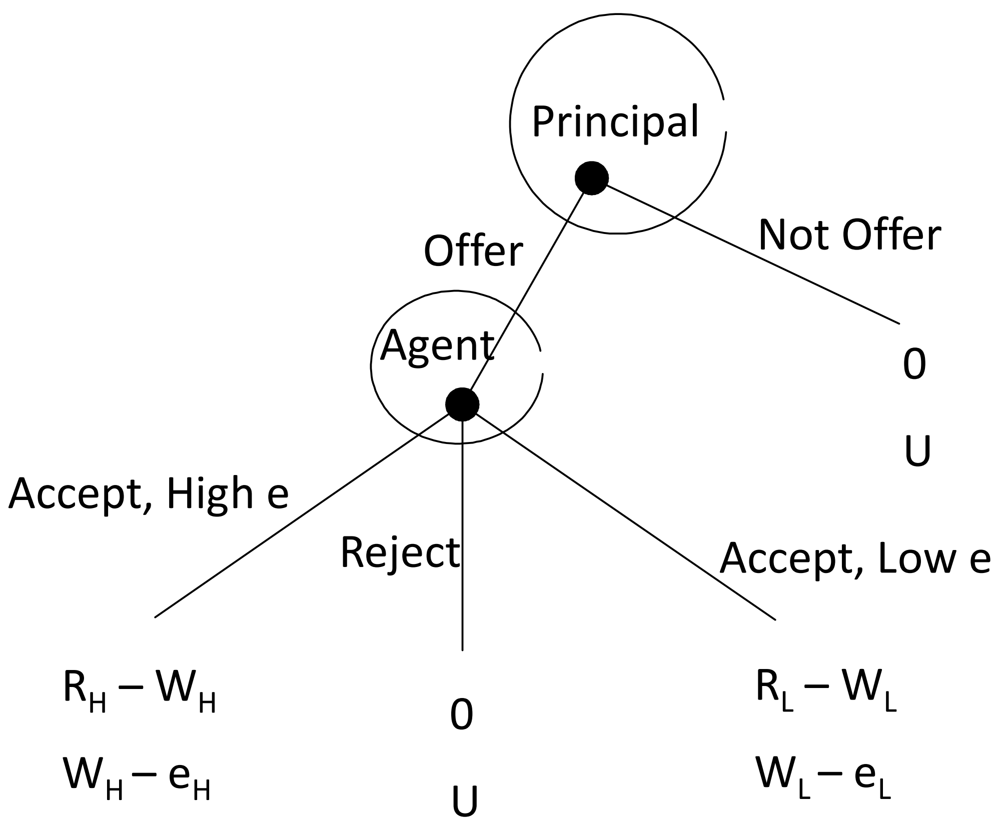

Figure 2: The game tree for the principal-agent problem with observable agent effort.

the agent the incentive to act in the principal's best interest. In this example, choosing high effort over low effort would mean $W_H - e_H > W_L - e_L$.

Participation constraint: The principal must structure the contract such that participation by the agent is better than non-participation. In this example, $W_H - e_H > U$.

Now, what should the principal set W_H to be in this example? Assuming that $R_H - W_H > R_L - W_L$, the principal wants the agent to exert high effort. The principal needs $W_H - e_H > W_L - e_L$ and $W_H - e_H > U$. The principal can guarantee that the agent will exert high effort if the contract is accepted by setting W_L such that $U > W_L - e_L$. With that choice for W_L, the principal will satisfy the agent's incentive compatibility constraint if the participation constraint is satisfied because now:

$$W_H - e_H > U > W_L - e_L.$$

Thus the principal should set $W_H > U + e_H$. But for the principal to maximize profit the wage should be set as close to the boundary condition as possible, so $W_H = U + e_H + \varepsilon$, for some small $\varepsilon > 0$. If we set up the constraint as $W_H - e_H \geq U$ then we would get that the wage must equal the reservation utility plus the effort cost.

Note that the principal has set up the contract to maximize profit by choosing the lowest possible wage offer to the agent that will induce high effort on the part of the agent and by setting the wage offer that would induce low effort so low that the agent would not accept that contract. The agent is better off accepting this offer and exerting high effort than by rejecting the offer or by choosing to exert low effort.

Now that we have some restrictions on parameters and have let $W_H = U + e_H + \varepsilon$ and $W_L = U$ (or really any wage such that the agent chooses not to accept a low wage offer contract), we can solve this game using backward induction. The agent will choose to accept the offer and exert high effort because $W_H - e_H > U$ and $W_H - e_H > W_L - e_L$. The principal knows what the agent will do at the agent's decision node, so the principal knows that if a contract is offered the principal will receive $R_H - W_H$ and the principal will offer the contract because $R_H - W_H > 0$.

2.1.3 Unobservable effort

The case with unobservable effort is more complicated because now we would have to consider how the agent's effort affects the probability of each outcome. We will not discuss that case in detail as it only adds complicating features. However, there is a more mathematical derivation of both the observable effort case and unobservable effort case in the Appendix. One important result is that the principal should always be better off under the case of observable effort than unobservable effort; if not, the principal could simply offer the contract used when effort is unobservable to the agent even when effort is observable.

2.2 Auctions

When one thinks of an auction, the likely picture is that of an auctioneer standing at the front of a room receiving bids to auction off a rare item. Bidders are typically seen raising paddles as the auctioneer increases the bid or shouting out bids. However, there are other types of auctions that are used in a variety of settings to allocate goods and services. We will discuss some basic formats, equilibrium bidding strategies, and instances in which auctions are used in policy.

2.2.1 Auction formats

In this section I will describe the four basic auction formats that we will discuss. The description will include the process by which bids are submitted and the assignment rule for the winner. For now, consider only the cases where we have a single, indivisible unit for sale.

1^{st}-price sealed bid auction

Process All bidders submit a bid on a piece of paper to the auctioneer.

Assignment rule The highest bidder is awarded the object. The price that the high bidder pays is equal to his bid.

Examples Many procurement auctions are 1^{st}-price sealed bid. Procurement auctions are typically run by the government to auction off a construction job (such as paving a stretch of highway).

Descending Clock (Dutch) Auction

Process There is a countdown clock that starts at the top of the value distribution and counts backwards. Thus, the price comes down as seconds tick off the clock. When a bidder wishes to stop the auction he or she yells, "stop".

Assignment rule The bidder who called out stop wins the auction, and the bidder pays the last price announced by the auctioneer.

Examples The Aalsmeer flower auction, in the Netherlands, is an example of this type of auction.

2^{nd}-price sealed bid auction

Process Bidders submit their bids on a piece of paper to the auctioneer.

Assignment rule The highest bidder wins, but the price that the highest bidder pays is equal to the 2^{nd} highest bid. Hence the term 2^{nd}-price auction.

Examples eBay is a kind of 2^{nd}-price auction. If you think about the very last seconds of an eBay auction (or if you consider that every person only submits one bid), think about what happens. You are sending in a bid. If you have the highest bid you will win. You will pay an amount equal to the 2^{nd} highest bid plus some small increment. Thus if you submit a bid of \$10 and the second highest bid is \$4, you pay \$4 plus whatever the minimum is (say it's quarter). So you would pay \$4.25.

Ascending clock auction

Process A clock starts at the bottom of the value distribution. As the clock ticks upward, the price of the item rises with the clock. This is truly supposed to be a continuous process, but it is very difficult to count continuously, so we will focus on one tick of the clock moving the price up one unit. The idea is that this is the smallest amount that anyone could possibly bid – that is how the ticks on the clock move the price up. All bidders are considered in the auction (either they are all standing or they all have their hands on a button – some mechanism to show that they are in). When the price reaches a level at which the bidder no longer wishes to purchase the object, the bidder drops out of the auction (sits down or releases the button). Bidders cannot reenter the auction. Eventually only two bidders will remain. When the next to last bidder drops out, the last bidder wins.

Assignment rule The winning bidder is the last bidder left in the auction. The bidder pays a price equal to the last price on the clock.

Examples The typical example given is Japanese fish markets, though those may be an urban legend. Thus, the English clock auction may only be a theoretical construct. However, the process is fairly similar to the standard portrayal of an auction with an auctioneer taking bids and asking for new high bids.

2.3 Bidding strategies

The previous section is meant to introduce you to the auction formats. In this section we will discuss the Nash equilibrium bidding strategies. We will derive the bid functions for some simple cases. For those truly interested in the details, I suggest reading Wolfstetter (1999).

2.3.1 General Environment

We discussed the assumptions of the general environment when discussing the revenue equivalence theorem – the Symmetric Independent Private Values environment with Risk-neutral bidders (SIPV-RN). These are the assumptions we are making here. Recall that a player's value, v_i, is drawn from some probability distribution. For simplicity, we will assume that all player values are drawn from the uniform distribution on the unit interval. This assumption means that all values are drawn from the interval $[0, 1]$ with equal probability. More importantly, if you draw a value of 0.7, then the probability that someone else drew a value less than you is also 0.7. We will not allow for the fact that someone else could draw the exact same value (theoretically, ties cannot occur with positive probability in a continuous probability distribution). Because probabilities must add up to 1, and because the other player's value draw must either be greater than your value or less than your value the probability that the other player has a value greater than yours is $1 - 0.7 = 0.3$.

The bidder's utility function is:

$$u(x) = \begin{cases} x\ if\ win\ the\ auction \\ 0\ if\ don't\ win \end{cases}$$

The term x in the utility function can typically that of as $v_i - p$, where v_i is bidder i's value and p is the price paid by bidder i. Note that a player's expected utility in these auctions can be noted as:

$$u_i = \Pr(win) * (v_i - p) + \Pr(lose) * 0$$

where $\Pr(win)$ is the probability that bidder i wins the auction and $\Pr(lose)$ is the probability that bidder i loses the auction. If the bidder wins he receives his value minus the price paid, or $(v_i - p)$ and if he loses he receives 0. Thus, for many auctions, the expected utility of a bidder is:

$$u_i = \Pr(win) * (v_i - p)$$

Note that the difficulty in deriving the theoretical results lies in establishing the probability of winning, $\Pr(win)$ and, in some cases, the price paid, p, particularly when the price paid depends on another bidder's bid.

In these auctions, a Nash equilibrium[7] is a bid function $b_i(v_i)$ that tells us what bid a bidder with value v_i would make.

2.3.2 Ascending clock auction – bidding strategy

Consider the following example. Assume that your value is 10. The clock begins at 0 and ticks upward: 0, 1, 2, 3, ..., 9, 10, 11, 12, 13, ... The question is, when should you drop out of the auction? Consider three possible cases:

1. The clock reaches 11:

   In this case you should drop out. While you increase your chances of winning the item by staying in, note that you will end up paying more than the item is worth to you. You can do better than this by dropping out of the auction and receiving a surplus of zero. So, as soon as the price on the clock exceeds your value you should drop out.

2. The clock is at some price less than 10:

   In this case you should remain in the auction. If you drop out you will receive 0 surplus. However, if you remain in the auction then you could win a positive surplus. If you drop out before your value is

[7] Technically this equilibrium is a Bayes-Nash equilibrium.

reached you are essentially giving up the chance to earn a positive surplus. Since this positive surplus is greater than the 0 surplus you would receive if you dropped out, you should stay in the auction.

3. The clock is at 10:

   What happens when the price on the clock reaches your value? Well, if you win the auction you get 0 surplus and if you drop out you get 0 surplus, so regardless of what you do you get 0 surplus. We will say that you stay in at 10, and drop out at 11. For one thing, it makes the Nash equilibrium bidding strategy simple – stay in until your value is reached, then drop out. Another way to motivate this is to consider that peoples values are drawn from the range of numbers $[0.01, 1.01, 2.01, 3.01, ...]$ instead of $[0, 1, 2, 3, ...]$. However, assume the prices increase as $[0, 1, 2, 3, ...]$. It is clear that if you have a value of 3.01 you should be in at 3, while if you have a value of 3.01 you should be out at 4.

So what would be the Nash equilibrium strategy in this auction? Stay in until your value is reached and drop out as soon as it is passed by the clock. *Or, if* we let $b_i(v_i)$ represent player i's bid as a function of his value, we have $b_i(v_i) = v_i$.

2.3.3 2^{nd}-price sealed bid auction – bidding strategy

In this auction you submit a bid and pay a price equal to that of the 2^{nd} highest bid. How should you bid?

One method of finding a Nash equilibrium (or a solution in general) is to propose that a strategy is a Nash equilibrium and then verify it. Naturally, it is a good idea to propose the right strategy the first time. So, consider the strategy: submit your value. Is this a good strategy?

What else could we do? We could submit a bid greater than the value or less than the value. Let's examine each of these.

Bid above your value Suppose we submit a bid above our value. What could this possibly change? Well, if we were to win when submitting our value then absolutely nothing changes – we still pay the same price because the price (if we win) is not tied to our bid. What happens if we submit a bid greater than our value and this causes us to switch from losing the auction to winning the auction? Suppose our value is 12 and the other player's value is 14. The other player submits 14 and we submit 12. We lose and earn 0 surplus. Now suppose we were to bid 15. We win, which is good, but we have to pay 15 for something that is only worth 12 to us. So we earn a "surplus" of (-3). This is bad. We could have done better by placing a bid of 10 (our value) and earning 0. So placing a bid equal to our value is better than placing a bid above the value in this case.

Bid below your value Suppose we submit a bid below our value. What could this possibly change? Well, if we were going to lose by submitting our value, then we still lose when submitting a bid below the value. So this changes nothing (at least not for us – it would help the highest bidder if we were the 2^{nd} highest bid!) as we still receive 0 surplus. Suppose we lower our bid and still win – again nothing changes because the 2^{nd} highest bidder has still submitted the same bid. It is possible though that we lower our bid and lose – here's where the problem occurs. Suppose our value is 12 and the other value is 11. We submit a bid of 12, we win, and we get a surplus of $(12 - 11) = 1$. Now suppose we submit a bid of 8 – we go from getting a surplus of 1 to getting a surplus of 0. It would be much better to submit a bid equal to your value and get a surplus of 1.

To further illustrate the point consider the following table when there are two bidders. Suppose that bidder 1 has a value of 12.

	Bidder 1's bid ($v_1 = 12$)		
Other bidder's bid	$b_1 = 10$	$b_1 = 12$	$b_1 = 14$
$b_2 < 10$	$v_1 - b_2$	$v_1 - b_2$	$v_1 - b_2$
$10 < b_2 < 12$	0	$v_1 - b_2$	$v_1 - b_2$
$12 < b_2 < 14$	0	0	$(v_1 - b_2)$
$b_2 > 14$	0	0	0

Note that $(v_1 - b_2)$ is NEGATIVE when $12 < b_2 < 14$. We have now determined that submitting a bid equal to our value is at least as good as submitting a bid greater than or lower than the value in some cases, and strictly better in other cases. Therefore, submitting a bid equal to your value is a *weakly dominant strategy*. A strategy is weakly dominant if it always provides a payoff at least as high as any alternative strategy, regardless of what the other player does. Thus, the Nash equilibrium for a 2^{nd}-price auction: Submit a bid equal to your value, so again we have $b_i(v_i) = v_i$.

You should note that the 2^{nd}-price sealed bid auction and the ascending clock auction are *strategically equivalent*, which means that all players have the same bidding strategies in either auction, even though the mechanism that produces the winner of the auction is slightly different.

2.3.4 1^{st}-price sealed bid auction – bidding strategy

In this auction you pay an amount equal to your bid if you win. The first question is, should you submit a bid equal to your value?

Bid equal to your value If you submit a bid equal to your value then you will expect to earn 0 surplus. If you win, then you will have to pay an amount equal to your value and if you lose you receive nothing. It stands to reason that you may be able to do better than this by submitting a bid below your value. The question is how far below your value?

Bid equal to the lowest possible value If you submit a bid equal to the lowest possible value that could be drawn then you will also receive 0 surplus. The reason is that you will never win because your bid was so low. Taken together with the fact that you will bid below your value, this means your actual bid should fall between the lowest possible value and your value draw.

Actual problem The actual problem facing someone bidding in a 1^{st}-price sealed bid auction is to maximize their expected utility. We have seen that in general we have:

$$u_i = \Pr(win) * (v_i - p)$$

In the case of the 1^{st}-price sealed bid auction, we know that if the bidder wins he will end up paying a price equal to his bid, so $p = b_i$. Thus we have:

$$u_i = \Pr(win) * (v_i - b_i)$$

For the example we have discussed, with two bidders and values uniformly distributed on $[0, 1]$, the $\Pr(win) = b_i$. Thus, if you bid 0.7 then you have a 70% chance of winning the auction; if you bid 0.4 then you have a 40% chance of winning the auction.[8] The problem is now a maximization problem:[9]

$$\begin{aligned} \max_{b_i} u_i &= b_i * (v_i - b_i) \\ \frac{du_i}{db_i} &= v_i - 2b_i \\ 0 &= v_i - 2b_i \\ 2b_i &= v_i \\ b_i &= \frac{1}{2} v_i \end{aligned}$$

Thus, given the SIPV-RN environment with values distributed uniformly on $[0, 1]$ and two bidders we have that player i's bid function is $b_i(v_i) = \frac{1}{2} v_i$. We can extend this result to the case of N bidders fairly easily. Now you have to consider the fact that your bid has to be higher than $N - 1$ other bidders' bids. The addition of more bidders alters the expected utility function to:

$$u_i = (b_i)^{N-1} * (v_i - b_i)$$

[8] We are assuming that the bid function is strictly monotone increasing, meaning that bidders with higher values will submit strictly higher bids.

The actual determination of the equilibrium bidding strategies is more complicated because we would need to show that the bid functions for the two bidders are best responses to each other.

[9] With multiple types there is no simple way to show this result graphically.

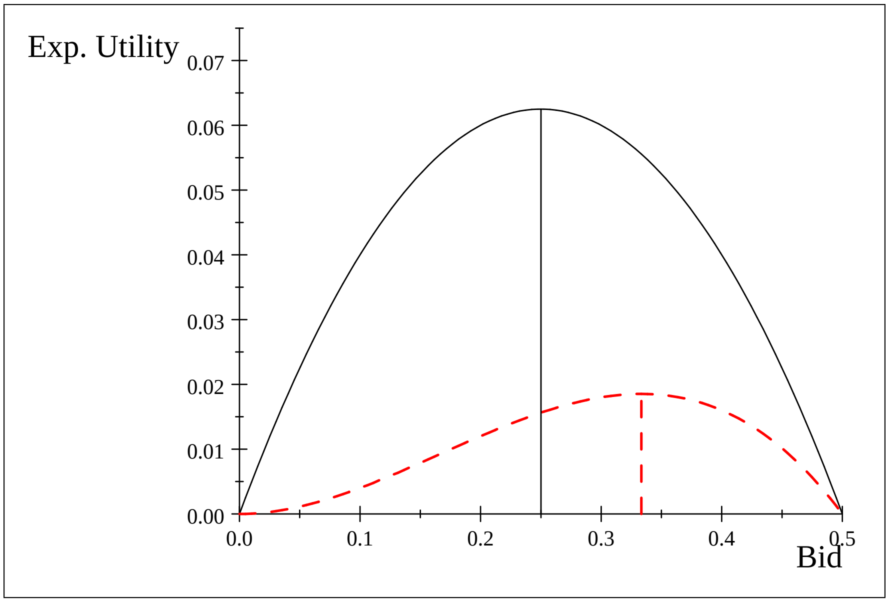

Figure 3: Expected surplus from potential bids in a 1^{st}-price sealed bid auction. The black curve assumes two bidders and the red curve assumes three bidders. The bidder has value $v_i = \frac{1}{2}$ and values are drawn from the uniform distribution on $[0, 1]$.

The reason that the b_i is raised to the $N - 1^{st}$ power is because the bidder now has to have a higher bid than $N - 1$ other bidders. While the problem is slightly more complicated than with 2 bidders it is still a fairly straightforward problem to solve:

$$
\begin{aligned}
\max_{b_i}\ u_i &= (b_i)^{N-1} * (v_i - b_i) \\
\frac{du_i}{db_i} &= (N-1)(b_i)^{N-2} * (v_i - b_i) + (b_i)^{N-1} * (-1) \\
0 &= (N-1)(b_i)^{N-2} * (v_i - b_i) + (b_i)^{N-1} * (-1) \\
0 &= (N-1)(b_i)^{N-2} * (v_i - b_i) - (b_i)^{N-1} \\
& \text{divide by } (b_i)^{N-2} \\
0 &= (N-1)(v_i - b_i) - b_i \\
0 &= (N-1)v_i - (N-1)b_i - b_i \\
0 &= (N-1)v_i - Nb_i \\
Nb_i &= (N-1)v_i \\
b_i &= \frac{N-1}{N} v_i
\end{aligned}
$$

Thus, for the general case of N bidders, the bid function in a 1^{st}-price sealed bid auction is $b_i(v_i) = \frac{N-1}{N} v_i$. Note that when $N = 2$ we have that $b_i(v_i) = \frac{1}{2} v_i$, which is what we found above. Thus you are shading your bid depending on how many other bidders there are. The more bidders, the less you shade your bid.[10]

Figure 3

shows the relationship between the expected utility of a bidder with value $v_i = \frac{1}{2}$ from making any bid between 0 and $\frac{1}{2}$ for the cases of two total bidders (black curved solid line) and three total bidders (red curved dashed line). The equilibrium bid when there are two bidders would be $b_i = \frac{1}{2} * v_i = \frac{1}{2} * \frac{1}{2} = \frac{1}{4}$ and when there are three bidders would be $b_i = \frac{2}{3} * v_i = \frac{2}{3} * \frac{1}{2} = \frac{1}{3}$. Note that Figure 3 shows that the maximum expected utility is attained when those bids are made for those respective cases.

[10] I will not ask you to derive results like these but it is instructive to know how the bid functions are determined.

2.3.5 Dutch auction – bidding strategy

Recall that with a Dutch auction the bidder watches as the clock descends, and then calls out when he sees a price that he wishes to pay. The problem facing the bidder is to maximize their expected utility. Their expected utility can be written as:

$$\begin{aligned} u_i &= \Pr(win) * (v_i - b_i) + \Pr(lose) * 0 \\ u_i &= \Pr(win) * (v_i - b_i) \end{aligned}$$

Notice that this is the same problem faced in the 1^{st}-price auction, if we make all the same assumptions we did when deriving the bid function for the 1^{st}-price sealed bid auction. Because the maximization problems are the same the 1^{st}-price and the Dutch auctions are strategically equivalent. Thus, the bidding strategy in the Dutch auction is to yell out stop when the clock reaches $\frac{N-1}{N}$ of your value, where N is the number of total bidders in the auction (including yourself)

2.4 Which format is "better"?

Now that we have seen the different formats, the question turns to which one is better. Better can mean many things, but we will focus on two meanings of better. From the standpoint of a benevolent social planner, better could mean more efficient. We will say that an auction is efficient if the item goes to the person with the highest value. Of course, an individual seller does not necessarily care about social goals such as efficiency, but about the revenue that the auction will generate for himself. The relevant question for the individual seller is then which format generates more revenue. We will look at both of these notions of "better".

Importantly, the bidding strategy for the 2^{nd}-price auction is a truth-telling bidding strategy as the bidders reveal their values; for the 1^{st}-price auction, bidders do not submit bids equal to their their values. While both are direct mechanisms, only the 2^{nd}-price auction is truth-telling.

2.4.1 Efficiency

We will define the level of efficiency in an auction as $\frac{V_w}{V_H}$, where V_w is the value of the winning bidder and V_H is the value of the high bidder. Note that if the winner is the high bidder, then efficiency is 1 or 100%. The question is, in all of our auction formats will the bidder with the highest value bid more than, less than, or an amount equal to bidders with lower values? It is easy to see that in an ascending clock or 2^{nd}-price sealed bid auction that higher values lead to higher bids because bidders simply submit their values as bids. In the Dutch and 1^{st}-price auctions, the bid function is $b_i = \frac{N-1}{N} v_i$. The question is, who will submit the highest bid? It should be fairly easy to see that higher values will submit higher bids. Technically, we can say that the bid function is increasing in the value draw – as the value draw increases, the bid increases. Thus, bidders with higher values will submit higher bids, and the bidder with the highest value will submit the highest bid. These auctions will also be 100% efficient, assuming that all of our conditions hold and bidders use the Nash equilibrium bidding strategies. Thus, theoretically there is no difference between the efficiency of the 1^{st}-price sealed bid auction, the Dutch auction, the 2^{nd}-price sealed bid auction, or the ascending clock auction

2.4.2 Revenue

As for revenue, we know that the 1^{st}-price and Dutch auctions are strategically equivalent and that the ascending clock auction and the 2^{nd}-price are strategically equivalent. Thus we know that the revenue from the 1^{st}-price and Dutch will be equal and the revenue from the ascending clock auction and the 2^{nd}-price will be equal. From our earlier discussion, we also know that the revenue equivalence theorem holds given our bidding environment. Thus, we know that the revenue generated from these auctions is equal, at least in expectation.[11][12]

[11] In practice the revenue could differ even if the bidders draw identical bids in the two auction formats.

[12] There is a more formal derivation showing the revenue equivalence of the 1^{st}-price and 2^{nd}-price auctions in the Appendix.

2.5 Breaking Revenue Equivalence and Efficiency

If all formats are perfectly efficient and generate the same revenue in expectation, why do auctioneers prefer one type or the other? The assumptions we have made may not hold and if they do not hold then there is no guarantee the results hold.

2.5.1 Breaking Revenue Equivalence

Suppose that instead of risk-neutral agents we had risk-averse bidders. They still have the exact same problem as before – they want to maximize their expected surplus. In the 2^{nd}-price and ascending clock auctions, there was no "maximization" problem – bidders simply submitted their bids or dropped out when the clock reached their value. Thus, the strategy should not change in these types of auctions if bidders are risk averse since they can do no better following another strategy. Because the strategy does not change the expected revenue from the 2^{nd}-price auction is still the same.

Consider the 1^{st}-price auction. Bidders wanted to maximize their expected utility, given by:

$$u_i = b_i * (v_i - b_i)$$

However, in the risk averse case with 2 bidders they now have risk aversion. Assuming that their risk aversion takes the form of $u_i(x_i) = \sqrt{x_i}$ we now have:

$$u_i = b_i * \sqrt{(v_i - b_i)}$$

This maximization problem is slightly more tedious to work through, but still tractable:

$$\begin{aligned}
\max_{b_i} \ u_i &= b_i * (v_i - b_i)^{1/2} \\
\frac{du_i}{db_i} &= (v_i - b_i)^{1/2} + b_i \left(\frac{1}{2}\right)(v_i - b_i)^{-\frac{1}{2}}(-1) \\
0 &= (v_i - b_i)^{1/2} + b_i \left(\frac{1}{2}\right)(v_i - b_i)^{-\frac{1}{2}}(-1) \\
0 &= (v_i - b_i)^{1/2} - b_i \left(\frac{1}{2}\right)(v_i - b_i)^{-\frac{1}{2}} \\
b_i \left(\frac{1}{2}\right)(v_i - b_i)^{-\frac{1}{2}} &= (v_i - b_i)^{1/2} \\
\frac{1}{2} b_i &= v_i - b_i \\
\frac{3}{2} b_i &= v_i \\
b_i &= \frac{2}{3} v_i
\end{aligned}$$

Recall that when we had 2 risk-neutral bidders in the 1^{st}-price sealed bid auction each bidder used $b_i(v_i) = \frac{1}{2}v_i$. With 2 risk averse bidders who have $u_i(x_i) = \sqrt{x_i}$, each bidder will bid $\frac{2}{3}$ of their value. Thus, we can see that the bidder is going to bid more in the risk averse case. Intuitively, if the bidder were to bid $\frac{1}{2}$ of his value in the risk averse case the marginal benefit from increasing the bid (the increase in the probability of winning) would be greater than the marginal cost (the amount of surplus lost). So we increase the bid until the marginal benefit of increasing the bid equals the marginal cost, just like we do with many other applications in economics. Because the risk averse bidders are now bidding more in the 1^{st}-price auctions than in the 2^{nd}-price auctions, the 1^{st}-price auction generates more expected revenue than the 2^{nd}-price auction.

2.5.2 Breaking Efficiency

Suppose we want to "break" efficiency. The true version of the ascending clock auction has the price moving up continuously with the tick of the clock. However, we know that people do not have continuous values, or, even if they do, there is some rational minimum amount by which their values must increase. In the US the smallest value one can have for a good is a penny, so it is not a stretch to think that the smallest unit in which values can be denominated is a penny. If this is a case, then a clock which moves at the rate of 1 penny per second (or 1 penny per hour or 1 penny per half-second – the rate is not important, but the units that it counts are) will still be perfectly efficient in the sense that the highest valued bidder will get the object. However, consider a clock that increases the price at a rate of 1 penny per second. Now consider the following prices and the corresponding amount of time it will take to auction off objects of these values:

$10 – 16.67 minutes

$1 *million* – 3.17 years

$1 *billion* – 3170 years

It does not really seem "efficient" to take 3170 years to auction off an item. In fact, it seems quite inefficient. So what auctioneers will typically do is impose a minimum bid increment. This minimum bid increment is the minimum amount by which the clock will increase (or the minimum amount by which bidders must increase the bid if they wish to place a new bid). While this speeds up the process, the introduction of the minimum increment can also affect the efficiency results of auctions. For instance, suppose 2 players have values of $14.08 and $14.92 respectively. If the clock ticks up at $1 per second, then both bidders will drop out at $14. In this case, a tie is declared and we must use a tie-breaking mechanism. The tie-breaking mechanism is usually a coin flip or some other equal probability game. Thus, on average, the bidder with a value of $14.08 will get the item half of the time. As you can see, the minimum increment introduces the possibility of inefficiency into the auction process.

Additionally, if bidders in the 1^{st}-price or Dutch auction have different risk preferences, there is no guarantee the highest valued user will bid the highest amount because individuals with different risk preferences will shade their bids by different amounts. Under a truth-telling mechanism like the 2^{nd}-price sealed bid or ascending auction, risk preferences should not affect efficiency if the bidder's understand how to bid in the auction.

2.5.3 Common Value Auctions

Suppose that I am auctioning off a jar of coins. The jar is see-through, so that you all can see there are coins in the jar. I tell you they are all U.S. coins from 1965-present (prior to 1965 some U.S. coins, notably dimes and quarters, are made of silver and are worth more than their monetary denomination) and you can see various coins (pennies, nickels, dimes, and quarters) in the jar. However, no one is allowed to look insider the jar or take the coins out of the jar. I conduct a 1^{st}-price sealed bid auction for the jar of coins, where the winner gets the coins. Clearly, the monetary amount that each individual would receive is the same because the coins do not depend on the winning bidder. Bidders may have different utility for the coins because perhaps they do not want to carry around coins to spend, but let us assume all bidders will happily take money in coin form.[13] Alternatively, we can consider that the bidders have no disutility from the monetary unit and only care about the value of the money. How are individuals' values formed for this jar of coins?

This auction is different than the ones we have discussed previously. In the prior auctions bidders had different values for the same item. It is fairly simple to motivate that example - there are plenty of goods for which you and your friends would pay different amounts. However, in the jar of coins example, all bidders have the same value for the same item, but they likely have different estimates of the item's value. They will not know if their estimate is correct unless they win the item and take possession of it. The jar of coins example seems a bit contrived – after all, who would auction off a jar of coins? But there are plenty of examples that fit this particular type of value determination. Consider a seller who has discovered that there is oil on his property. The seller wants to sell because he does not know much about extracting oil and refining it, but the bidders will not know exactly how much oil is in the deposit until they own the property and can begin to extract it. Or perhaps a target firm is for sale. Other firms would like to buy this target firm, and they have an estimate of the target firm's value based on observable information, but they will not truly know the target firm's value until it is acquired.[14]

[13] It is also possible that the utility value differs between individuals because someone is a coin collector who might value a particular coin at more than its monetary value because it fills a whole in a collection.

[14] The genre of televsion shows about bidding for storage units also fits this model.

Auctions of this type are known as "common value auctions." They are different than "private value auctions" because bidders now have a signal about the items value, but do not know the true value until it is purchased. Making the concept slightly more formal:[15]

1. There is a common value V for the item, which is drawn from some underlying probability distribution.
2. Bidders receive a signal S_i of V prior to bidding. The signal S_i is drawn from some probability distribution $(V-\varepsilon, V+\varepsilon)$. Thus, each bidder's signal is dependent on the common value V, but they all have (potentially) different signals.

We will not dive deeply into the determination of this equilibrium, but will consider an example assuming a 1^{st}-price sealed bid auction. If we assume a symmetric equilibrium (meaning that all bidders use the same strategy), then the bidder with the highest signal will win because that bidder will bid the most. However, bidders who win know that they have the highest signal and, the more bidders in the auction or the larger ε is, the more likely that signal is an overestimate of V. Thus, they need to shade their bid not only to receive a surplus (as they would in a private value 1^{st}-price sealed bid auction) but also because they realize their signal, if they win, is likely an overestimate of V. In equilibrium, bidders in a common value auction make positive profit. However, bidders who are not familiar with the common value auction may (likely) end up bidding too much for the item, and ultimately lose money because they pay some price $P > V$. Overbidding in a common value auction and losing money is know as the *winner's curse*. The term really applies to common value auctions, though some individuals use it (likely incorrectly) with private value auctions. In a private value auction individuals know their values, and can avoid the winner's curse by making sure they do not submit bids that would lead to payments greater than their value, whereas in common value auctions bidders will likely be submitting a bid below their signal, but that bid might still be above the common value V. If a bidder wins a private value auction and feels like they overspent that would more accurately be described as *buyer's remorse*.

2.5.4 Auctions in Policy

Where are auctions used by policy makers? Earlier in the course we briefly mentioned that some regions use a cap and trade model to ration permits for carbon emissions. Those permits are initially auctioned. Also, in the U.S. the Federal Communications Commission (FCC) auctions off the rights to wireless spectrum. Another example, which we will not discuss in detail because it is essentially a 1^{st}-price sealed bid auction in which the lowest bidder wins the right to perform a job, is procurement auctions, also known as reverse auctions.[16]

Carbon Emissions (Cap and Trade) In our discussion of negative externalities it was mentioned that regulators may want to provide incentives for producers to reduce carbon emissions. There are measures that some may consider punitive that could be imposed – taxes, fines, and audits (or regulatory reviews) are some of those measures. While blatant acts of environmental harm may call for such punitive measures, reframing the incentives in a more positive light may yield better compliance.[17] A cap and trade system for carbon emissions is one such system. Under this system, greenhouse gas emitters purchase pollution permits at auction which allow them a certain level of emissions.[18] The providers are incentivized to reduce emissions because they can typically bank any unused permits for later use or sell any unused permits to other producers.[19] The overall cap is generally reduced over time to reduce emissions. One potential unintended effect of this process, noted in Jones (2016), is that the permit purchasers may be incentivized

[15] See Wolfstetter (1999), pgs. 226-229 for more information.

[16] New York State has information on the bidding process for contracting with the state: https://ogs.ny.gov/procurement/bidding-101

[17] Before the Fall 2020 semester began, most University policies concerning behavior during the pandemic focused on punishments (suspension, dismissal, etc.) I recommended a more positive approach might lead to better compliance and a small group met to discuss those possibilities. Ultimately, because very few classes were held in person, there were not many additional incentives, positive or negative, provided.

[18] https://www.edf.org/climate/how-cap-and-trade-works

[19] While there may be more, at least one alum of the PPOL program, Brian Jones, wrote his dissertation on the impacts of cap and trade (Jones, 2016).

to keep emissions near the overall cap so that the reduction in emissions over time happens more gradually. The California Air Resources Board, Regional Greenhouse Gas Initiative (RGGI), and Congressional Budget Office all provide additional information on processes and results.[20]

Spectrum Auctions The Federal Communications Commission (FCC) has been auctioning off the rights to wireless spectrum since the mid-1990s.[21] A recent policy goal of the FCC is to provide internet access to rural communities,[22] and they have been conducting auctions of rural broadband since 2018.[23] Rather than auctioning off slices of wireless spectrum as single goods, these auctions tend to be conducted as package, multi-unit, or combinatorial auctions.[24] The synergistic nature of wireless spectrum is the reason for these package type auctions – a purchaser needs access to a set of wireless spectrum slices in order to form a communication network and auctioning off individual slices without the opportunity for package bidding could lead to inefficiency or speculation.

Cramton (1997) provides an early assessment of the FCC spectrum auctions. The novelty is that the FCC used a simultaneous multiple-round auction which did not end until there was a round in which no bids were placed on any item. Cramton notes that increased competition in the final rounds suggests the auctions were efficient, but also notes that bidder alliances may have limited competition in some markets. Cramton and Schwartz (2000) examines how open bidding could lead to collusion. The "trailing digits" practice, in which the license numbers of interest to a particular bidder could be added as the last few digits of the bid, could be used to inform other bidders which licenses were of interest to a particular bidder and possibly to imply that bid retaliation would be used if those bidders continued to bid on that license.[25] They offer some suggestions on how to reduce potential collusive practices. Weber (1997) discusses how the multi-unit auction format may lead to strategic demand reduction, or bidders not placing bids even when there value is above the current bid for the item, in equilibrium and practice. Banks et al. (2003) discuss a combinatorial auction design as an alternative to the simultaneous multi-round auction used by the FCC and show that the combinatorial design outperforms the FCC design under some conditions. Plott and Salmon (2004) develop models of bidder behavior and apply them to the UK3B spectrum auction.

2.6 Matching

Throughout the course we have primarily used prices as our rationing mechanism, though at times "prices" may have been denominated in terms of other goods rather than currency. However, there are times when prices are not the relevant method for allocating resources. Consider a situation in which there are limited position openings and limited candidates for those openings. One example would be athletes in professional sports leagues. After players have attained a certain level of experience, a system of negotiations and prices is used to determine which team employs the player. In those cases, it is clear that the team and player have both expressed preferences and come to a mutual agreement that is beneficial to both. However, many professional sports leagues use a draft for new players entering the league. While that is a system of matching, it is mostly one-sided: the preference of the teams dominates because the only option for the players, if they want to play in that professional league, is to sign with the team that drafted them. They can choose not to sign or to sign with a different (perhaps foreign) league, but the preferences of those new players entering the league are generally not considered.[26] If a player is from Florida and really does not want to work in Chicago, the player does not have much leverage if drafted by a team in Chicago.

[20] https://ww2.arb.ca.gov/our-work/programs/cap-and-trade-program/auction-information
https://www.rggi.org/
https://www.cbo.gov/sites/default/files/111th-congress-2009-2010/reports/11-04-2010-cap-and-trade.pdf

[21] https://www.fcc.gov/auctions, fcc.gov/auctions-summary

[22] https://www.fcc.gov/about-fcc/fcc-initiatives/homework-gap-and-connectivity-divide

[23] https://www.fcc.gov/auctions/ruralbroadbandauctions

[24] There is a technical guide with details about the bidding process available at this link.
https://www.fcc.gov/auction/108/education

[25] One could ask why the bidder would increase their own payment in that way, but when bids are in the hundreds of thousands or millions of dollars, adding 234 as the last digits of a bid is an inconsequential amount to signal interest in license 234, particularly if that bid helps keep the final price of the license low.

[26] There are some examples of teams opting not to select certain players in drafts because they knew the player would not sign with the team, but those tend to be the exceptions. In the NFL, John Elway, Bo Jackson, and Eli Manning are three players who leveraged their ability and outside options to attain a match with a team that was more aligned with their preferences.

Matching algorithms can be used to better match preferences of job candidates and employers. One potential goal with matching algorithms is to use an algorithm that incentivizes both sides to truthfully reveal their preferences. If the results of the algorithm can be influenced by strategic behavior on either side, then the algorithm will be optimizing over the stated (potentially untruthful) preferences and the resulting matches are unlikely to be globally optimal, though all individuals may be behaving rationally to maximize their utility given what others are doing. One problem with matching algorithms is that they usually require a full set of preferences – or at least a full set of preferred options to not having a match. In practice, some matching processes only allow a limited number of potential matches to be specified, which opens up the possibility for strategic behavior. Both sides may be trying to guess at who wants to match with them and they may make their choices to ensure a match is made.

Another potential goal of matching algorithms is to reach a stable outcome, which is just that no pair would prefer to be matched to each other than with the partner they are currently matched. This goal matches our general concept of equilibrium – the system reaches a point at which there is no beneficial change. Roth (1982) proved (under certain assumptions) that there is no matching procedure in which both sides of the market reveal their true preferences and that yields a stable outcome, but there do exist matching procedures that lead to a stable outcome and one side truthfully revealing their preferences. If participants in the matching process are compelled to comply with the final match, attaining a stable outcome may not be the primary concern of the mechanism designer. Still, if too many participants are unsatisfied they may seek options outside the system.

Medical Residencies The National Resident Matching Program (NRMP)[27] matches medical students to open resident positions and fellowships. A matching program had existed since the 1920s, but was beset by strategic behavior on the part of hospitals (which made early offers due to a lack of supply of available residents) and medical students (which delayed accepting offers due to the lack of supply).[28] In the 1940s medical schools changed behavior, severely limiting the ability of residents to select a suitable match, and in 1952 the NRMP[29] was established "at the request of medical students to provide an orderly and fair mechanism for matching the preferences of applicants for U.S. residency positions with the preferences of residency program directors."[30] There were concerns that the original algorithm being used could be manipulated by students strategically, rather than truthfully, stating preferences, and that the algorithm favored the preferences of the programs rather than those of the students. That algorithm was updated in the late 1990s to favor applicant preferences over those of programs.[31] Roth (2003) provides a brief review of the NRMP while Roth and Peranson (1999) provide a more technical discussion.

School Choice School choice is another area in which matching mechanisms have been used. Roth (1985) is an early work in this area and relates to Roth (1982). The college admissions problem in the 1985 paper differs from the marriage problem in the 1982 paper because the matching is no longer one-to-one. While colleges and students have preferences for each individual match, colleges also need preferences over final allocations of students, making the matching problem more difficult. Abdulkadiroğlu and Sönmez (2003) is an early work that examines school choice mechanisms in Boston, Columbus, Minneapolis, and Seattle. They find that the plans in place at that time had shortcomings, primarily that the plans were confusing and required strategic choice by students/parents. Abdulkadiroğlu et al. (2005) focus specifically on the Boston school system and discuss two alternative mechanisms, deferred acceptance algorithms[32] and top trading cycles. Ergin and Sönmez (2006) examine the theoretical properties of the Boston mechanism, which matches as many students as possible to their top choice school. Thus, the students (parents) have

[27] https://www.nrmp.org/

[28] To tie back to the sports analogy, during this same time (1930s and 1940s) Major League Baseball teams bid against each other for young players. They too enacted rules to curb this behavior, first by creating a "Bonus Rule" that required the player to spend two years on the major league team and then, ultimately, by creating an amateur player draft.

[29] Here is a brief description of the proposal in the September 29, 1951 volume of the Journal of the American Medical Association:
https://jamanetwork.com/journals/jama/article-abstract/312075

[30] https://www.nrmp.org/about-nrmp/

[31] Here is a tutorial on how the NRMP matching algorithm works:
https://www.youtube.com/watch?v=kvgfgGmemdA

[32] Similar to the NRMP matching design in which seats are assigned tentatively.

an incentive to strategically choose their top school to be one from the acceptable set of schools at which they believe there will be an opening. They show that replacing the Boston mechanism to a student-optimal stable mechanism would lead to unambiguous welfare gains. Chen and Sönmez (2006) conduct a laboratory experiment to examine behavior in a controlled setting under the Boston mechanism, Gale-Shapley, and top trading cycles. Gale-Shapley and top trading cycles have superior theoretical properties and they find that participants list schools strategically under the Boston mechanism which leads to inefficient outcomes. Pais and Pintér (2008) also study those three mechanisms experimentally. They find that top trading cycles outperforms Gale-Shapley and the Boston mechanism in efficiency and that manipulation of reported preferences is strongest under the Boston mechanism. Both experimental papers recommend adopting an alternative mechanism than the Boston mechanism.

As new mechanisms are implemented their properties are more fully fleshed out as more realistic assumptions are made. Haeringer and Klijn (2009) examine the problems that arise when students are limited in the number of schools they can list, which is typically a practical consideration. They find that the Boston mechanism generates stable outcomes while the top trading cycles does not. Calsamiglia, Haeringer, and Klijn (2010) study this practical consideration experimentally and also highlight the role of the school district. Examining efficiency results, they find that top trading cycles outperforms student optimal stable matching (deferred acceptance) which outperforms the Boston mechanism. However, they find less encouraging results for stability.

Erdil and Ergin (2008) consider that schools (and possibly students) do not have strict preference orderings over students (possibly schools). Due to the lack of strict preferences on the part of the schools, the tie-breaking rule used to assign students to schools can cause efficiency losses. Abdulkadiroğlu, Che, and Yasuda (2011) reconsider the Boston mechanism and argue that the welfare gains by other mechanisms are made on unrealistic assumptions, though they do note that the Boston mechanism is not a truth-telling mechanism. Kesten (2010) examines an efficiency-adjusted deferred acceptance mechanism and finds that it can recover welfare losses induced by the student-optimal stable matching mechanism. More recently, Lien, Zheng, and Zhong (2016) and Dur, Hammond, and Kesten (2021) examine school choice mechanisms experimentally. Both study timing aspects of the submission process, with Lien et al. focusing on ex-ante fairness and Dur et al. focusing on sequential preference submission. Both find that the Boston mechanism compares favorably with the counterpart mechanisms in their study and suggest there may be important policy implications for those systems using the Boston mechanism.

As you can see, the literature has progressed from an initial model under a one-to-one matching system to one in which the matching is one-to-many. Theoretical models were constructed and comparisons of the theoretical properties were proved and tested in the laboratory. As the models were implemented practically, new challenges arose and those challenges were examined theoretically and experimentally.

2.7 Voting

We have considered mechanisms between individuals, between a single seller and multiple buyers, and between multiple agents on each side of a market. In each of these scenarios the mechanisms are used to elicit agents' preferences and make allocations. Depending on the mechanism, agents may report preferences truthfully or not. The goal of the mechanism designer is to create a mechanism that takes into consideration how the individuals will report preferences and uses that anticipated behavior to make efficient allocations. Another area in which the principles of mechanism design may be used to elicit preferences and establish a preferred outcome is voting. While the voting systems typically used in high profile political elections in the U.S. tend to take a similar form (either majority or plurality voting), these are not the only voting systems.

Although some may believe that voting should be done without strategic considerations, one could argue that voting is a game played between players where the outcome depends on the actions (votes) of all players. Thus we will analyze voting from a strategic perspective and compare outcomes of various voting mechanisms when people vote strategically. *Sincere voting* is voting for your preferred outcome without consideration of other voters' preferences or how you could impact the outcome of the election by changing your vote. In contrast, *strategic voting* considers the preferences of other voters as well as the individual's preferences over the remaining options and the impact the individual vote can have upon the outcome of the election. Note that voters may still find it in their best interest to vote for their preferred choice when voting strategically, it is just that they went through the process of determining that voting for their preferred choice would leave them best off after the election has taken place. Any Nash equilibrium of a voting system will require that voters cast their votes in a manner such that no voter can alter the outcome of the election by changing their vote to make themselves better off given what the other voters are doing. Thus, sincere voting is typically not a Nash equilibrium in most of these mechanisms that we will discuss.

2.7.1 Candidate Location

Before discussing voting behavior we begin with a discussion of how candidates might determine their positions on political issues. We examine the two and three candidate cases as there are important differences between the two.

Two Candidate Elections When there are only two candidates (or choices) and no later rounds of voting, there can be no strategic behavior on the part of the voters because voters who are not voting sincerely for their most preferred choice are, by construction of the election, voting for their least preferred choice. For a simple two candidate election, we will briefly examine the decisions of the candidates running in an election and assume that voters simply choose their preferred candidate.

We know that candidates represent many different issues in an election – however, voters are only able to cast one vote, so they must take all the available information as well as their own personal feelings (utility) about how strong a stance a candidate takes on particular issues and map it into a ranking for the candidates. We will assume that voters rank candidates along the "political spectrum". The political spectrum runs on the unit interval (from 0 to 1), with 0 representing the extreme left and 1 representing the extreme right. Candidates will choose a location, L_i, on the political spectrum. A location of 0.5 means the candidate is directly in the center of the political spectrum; a location of 0.6 means the candidate is right of center and a location of 0.4 means the candidate is left of center. The two candidates will choose their locations simultaneously. As there are only two candidates we will use a simply majority rule voting system and if the candidates receive the same amount of votes then there will be a coin flip to determine the winner.. We will assume that voters are distributed uniformly along the unit interval – essentially there is one voter located at every spot along the unit interval.[33] Voters will vote for the candidate who is nearest to their own preferences. If both candidates are the same distance away from a voter then the voter flips a coin to choose between the two candidates. Where should the candidates locate?

Suppose they begin by locating at the extreme points along the political spectrum. Both receive 50% of the vote. However, each candidate has an incentive to change location. Consider the candidate who located at 0 (call this Candidate A). Candidate A now has the incentive to move just to the left of Candidate B (location around 0.999) to take all but one vote in the election. However, if candidate A does this then candidate B has the incentive to locate just to A's left, at 0.9998. Both candidates continue this move leftward, until one of them reaches the location at 0.5. Now, if the other candidate moves a little to the left then the candidate still loses. That candidate's only hope is to locate at the exact same spot as the other candidate, at 0.5. Now both candidates are at 0.5. Does either have an incentive to move? No, neither can obtain a larger vote share by choosing another location. Thus we have found a Nash equilibrium to the game: both candidates locate at the midpoint.

You should note that the candidates locating at an identical location that is not 0.5 is NOT a Nash equilibrium to this game. For example, if both candidates locate at 0.75 then they both have an incentive to move to 0.74 because they will gain more votes at that position. This result about candidate locations on the political spectrum is known as the Median Voter Theorem. More formally, suppose we have two candidates and voters have single-peaked preferences (meaning they only have one peak along the political spectrum, or, if you exclude the boundary points of 0 and 1, their preferences exhibit no local minima along the political spectrum) along the political spectrum. If we let m^* be the location of the median voter along the political spectrum, then the Nash equilibrium location choices by the two candidates are (m^*, m^*).

In practice there are a number of reasons why candidates may not locate exactly in the middle of the political spectrum. One is that there may be more than two candidates. Another is that the assumptions of the model may be violated. A third is that there are multiple issues under consideration during an election

[33] Uniformly distributed voter preferences is a simplifying assumption. The same general result holds if we have voter preferences that are distributed symmetrically around the midpoint of the political spectrum.

and candidates may take more extreme stances on different issues that better align with their political party preferences.

Three Candidate Elections Now, consider an election with three candidates and the same setup as the two candidate election. The candidates are to choose their location along the political spectrum described. We can show that there is no Nash equilibrium that does not involve probabilistically choosing locations[34] if their goal is to maximize vote share.[35] Although the proof is a little involved, the basic idea is if they all locate at the same spot then someone has an incentive to shift to the side with the most mass. If they all locate at different spots then the two candidates on the outside have an incentive to move inward, and at some point as the candidates move inward it will be in the best interest of the candidate in the middle to move from being between the other two candidates to being just on the outside of the candidate that has the most mass on his side. While the location decisions of the candidates in a three candidate election is interesting, our goal is to understand how having three candidates in an election allows for strategic voting.

1992 US Presidential Race - example In the 1992 U.S. Presidential race the three major candidates were Clinton, Bush, and Perot. Given their final vote shares (modified very slightly to fit the example), if we were to locate our candidates along the political spectrum with uniformly distributed voter preferences they would locate at 0.25, 0.62, and 0.99, for Clinton, Bush, and Perot respectively.[36] These locations would give Clinton a 43% share, Bush a 37% share, and Perot a 20% share of the popular vote if voters vote sincerely. We will assume that the plurality winner is elected, abstracting away from the real-world complicating factor of the electoral college.

Given these locations we can back out what the preferences of the voters must be assuming they all voted sincerely. The preference rankings, along with the percentage of the population that have those preferences, are included in the table below. Note that these preferences and rankings pertain to how the example is structured, and is not based on any survey of actual individuals.

Percentage	43% (CBP)	20% (PBC)	19% (BCP)	18% (BPC)
1^{st}	Clinton	Perot	Bush	Bush
2^{nd}	Bush	Bush	Clinton	Perot
3^{rd}	Perot	Clinton	Perot	Clinton

I found these preferences and percentages as follows. I assumed anyone who sincerely voted for a candidate must rank that candidate first. I assumed anyone who ranked Clinton first must prefer Bush to Perot because Bush is more left than Perot. I assumed anyone who ranked Perot first must prefer Bush to Clinton because Bush is more right than Clinton. People who voted for Bush could rank Clinton or Perot 2^{nd}, depending on which candidate is closer to them. If we remove Bush from the election we find that Clinton defeats Perot, and the vote share for the candidates is 62% to 38%. Thus, Clinton picks up 19% and Perot picks up 18%, which gives us the rankings above.

Now, what would happen if these blocks of people voted strategically? It all really comes down to those who favor Perot – they favor Clinton the least, and by voting for Perot they allow Clinton to win. Thus, if they voted strategically they could alter the election and end up with their second-best choice if they vote for Bush. Strategic voting by the PBC voters would give Bush a 57% vote share to Clinton' 43% vote share. You should notice that no other group has enough power to swing the election. If the 18% BPC people attempted to vote strategically, then Perot would have 38% of the vote to Clinton's 43% and Bush's 19%. However, neither the CBP people nor the BCP people would wish to change their vote to Perot. While I have constructed this example based on very simple assumptions, it illustrates how strategic voting could affect elections.[37]

[34] We have not discussed games like Rock, Paper, Scissors in which there is no "pure strategy" Nash equilibrium to the game where both players choose a particular strategy with certainty but in which the Nash equilibrium of the game is to randomize over the possible strategies.

[35] If the goal is to maximize the probability of winning then there is an unusual Nash equilibrium where one candidate locates at the midpoint and the other two locate just to the left and right of the midpoint. The two candidates to the left and right end up tying for the most votes and winning with 50% probability, but no candidates can increase their odds of winning by choosing a different location.

[36] These were almost certainly not the locations of the candidates in the actual election, but these are the candidate locations that would be backed out under our assumptions.

[37] See Burden (2005) for a discussion about the effect of minor parties on strategic voting in U.S. presidential elections.

2.7.2 Alternatives to Majority/Plurality Voting

Majority and plurality voting are likely the most recognizable voting methods, in part because of the ease in which a winner is determined. However, there are alternatives to majority/plurality voting. These alternative systems have their own benefits and drawbacks.

Condorcet Election A Condorcet election is an election in which candidates are paired in a round-robin type of voting tournament. For instance, in the three candidate U.S. presidential election discussed earlier we would have three separate elections: one for Bush vs. Perot, one for Bush vs. Clinton, and one for Clinton vs. Perot. A candidate who goes undefeated in this type of election is called the Condorcet winner. If we hold a Condorcet election using the preferences above, and people vote sincerely, we find that Bush beats both Perot (80% to 20%) and Clinton (57% to 43%). This result means that Bush is the Condorcet winner and is elected. Already we can see the tension building between the different voting systems – plurality gives us Clinton as a winner and Condorcet gives us Bush. It is elections in which these types of results could occur in which people may question whether the "right" voting mechanism is being used.

Now suppose that we have a committee of three attempting to choose a flavor of ice cream for their next organizational meeting. The committee members' preferences are as follows:

Person	Greg	Peter	Bobby
1^{st}	Chocolate	Vanilla	Strawberry
2^{nd}	Vanilla	Strawberry	Chocolate
3^{rd}	Strawberry	Chocolate	Vanilla

They have decided to use a plurality rule and have not chosen a tie-breaking mechanism. If individuals vote sincerely, all three flavors receive a single vote. Now, any of these three individuals could vote strategically to break the tie – Greg could vote for vanilla and guarantee that vanilla wins if the others vote sincerely, so at least Greg would be guaranteed to receive his second favorite ice cream.

Suppose they decide to use a Condorcet election. With chocolate vs. vanilla, chocolate wins. With vanilla vs. strawberry, vanilla wins. One would think that if chocolate defeats vanilla, and vanilla defeats strawberry, that chocolate would beat strawberry. However, strawberry beats chocolate in the actual vote, and the committee members are left as they were before, with each flavor ending up defeating one flavor and being defeated by another flavor. Thus there does not necessarily need to be a Condorcet winner in a Condorcet election.

Sequential Pairwise Voting Suppose that Greg suggests a voting procedure that will ensure a winner. Instead of a Condorcet election, there will be a series of sequential single-elimination elections. The committee members will take two of the flavors and vote on them in a first round, with the winner being matched against the remaining flavor. Because Greg suggested the voting procedure the committee members agree to allow chocolate to have the bye into the second round. In the first round, if the members vote sincerely, vanilla beats strawberry. Then, when vanilla is matched against chocolate the winner is chocolate. Finally, they have a flavor of ice cream for their next meeting and they can move on to other issues.

However, we have assumed sincere voting. First, chocolate is Peter's least favorite ice cream. He knows that if vanilla (his favorite) is going to go up against chocolate it will lose. So if Peter acts strategically he can change his vote to strawberry in the first round, which is his second favorite flavor, because we know that strawberry will beat chocolate when matched against each other. Thus, strategic voting can run rampant in all rounds of a sequential voting system except for the last as in the last round it is essentially a two candidate election with no future rounds, so there is no need for strategic voting.

Second, the order of the elections is extremely important. If Greg believes that everyone votes sincerely then he wants chocolate to "have the bye" into the last round because chocolate will beat vanilla. However, if Greg believes that the members will vote strategically, then he would suggest that vanilla has the bye into the last round and that chocolate should be matched against strawberry in the first round. If this occurs then Greg would vote for chocolate in the first round, and Michael, knowing that his preferred outcome of strawberry will lose to vanilla in the second round, will vote for chocolate in the first round. Thus Greg can appear to be "nice" by allowing another's favorite to receive the bye into the last round, while in reality he is acting strategically so that his favorite flavor will be chosen.

Borda Counts In many election systems only the top choice of the voters receives any consideration. It could be the case that in a plurality election 43% of the people have one candidate as their favorite but the other 57% despise that candidate and rank that candidate last. However, because there are two more candidates who divide the remaining 57% of the vote the candidate who is despised by 57% of the population ends up winning. Borda counts were devised to give weight to more than one choice among voters' preference rankings. Under a Borda count method, voters are asked to rank candidates in slots 1 through N, where N is either some predetermined cut-off number or determined by the number of candidates. A weight is assigned to each position in the ranking order, and total points are calculated based on how the candidates are ranked and the assigned weights. The candidate with the most points is declared the winner.

Borda counts are frequently used in sports. The NCAA polls use Borda counts and many awards in the major professional sports are also based on Borda counts. However, as we have seen in all of these voting systems, there are potential problems. The first problem is that strategic voting can run rampant as there is no guarantee that voters rank the candidates sincerely. Thus, if all the candidates are to be ranked, then voters can attempt to lower the point totals of candidates they think will compete with their favorite by ranking the candidate lower. Or, if a voter's preferred candidate will not win, then voters have the incentive to put a lower ranked choice in their top spot to give that choice more points. Also, if the number of candidates exceeds the cut-off number then voters may just leave candidates who will compete with their preferred candidate completely off the ballot, awarding that candidate no points.

A second consideration is the weights assigned to the rankings. The NCAA polls use a constant decrement of 1 point between their rankings. A first-place vote receives 25 points, a second-place vote receives 24 points, ..., down to a twenty-fifth place vote receiving 1 point. In Major League Baseball Most Valuable Player award voting, a first-place vote receives 14 points, while a second-place vote receives 9 points. The increments then decline at a rate of one less point for one lower spot in the ranking, so that a tenth place vote receives 1 point. There is no good answer as to which weighting system is "correct" as the weighting varies depending on how much more benefit a candidate should receive for being listed first rather than second.[38]

Multi-seat Elections The election procedures we have discussed have focused on determining a single election winner. However, in some elections, such as those for school boards, there is not a single winner but multiple winners. As with single-seat elections, there are multiple methods of determining winners in multi-seat elections. While we will not discuss those methods, you should be aware that, like with single-seat systems, multi-seat election methods have benefits and drawbacks.

3 Appendix

3.1 Principal-Agent (more formal discussion)

Consider a similar structure as the prior examples with a principal and an agent. Now, however, the principal's profit is continuous over the range $[\underline{\pi}, \overline{\pi}]$. The agent's effort $e \in E$ can be any effort level from the set of effort levels. The agent's effort choice affects the profit stochastically – thus any profit can occur with any effort choice. The conditional density function $f(\pi|e)$ describes the relationship between profit and effort, and $f(\pi|e) > 0$ for all $e \in E$ and all $\pi \in [\underline{\pi}, \overline{\pi}]$. We again restrict the agent's effort choice to $\{e_H, e_L\}$ but assume that the distribution of π conditional on e_H first-order stochastically dominates the distribution conditional on e_L. Thus, the expected level of profits when the agent chooses e_H is greater than when he chooses e_L:

$$\int_{\underline{\pi}}^{\overline{\pi}} \pi f(\pi|e_H)\, d\pi > \int_{\underline{\pi}}^{\overline{\pi}} \pi f(\pi|e_L)\, d\pi.$$

The agent has utility over wage and effort, $u(w,e)$. We focus on a slightly more general case than above, with $u(w,e) = v(w) - g(e)$. Assume that $v'(w) > 0$, $v''(w) < 0$, and $g(e_H) > g(e_L)$. The principal receives the profit realization π but must pay the wage from that profit, so ultimately the principal receives $\pi - w$.

[38] Benoit (1992) examines how changing the weights in MLB MVP voting could alter the outcome. Of the 86 elections Benoit examined, 24 of them could have had a different winner if a different point awards system was used.

3.1.1 Observable effort

Again, consider the case of observable effort. The principal offers a contract that the agent can accept or reject. The contract specifies the effort level $e \in \{e_L, e_H\}$ and the wage as a function of observed profit, $w(\pi)$. The agent's reservation utility of accepting the contract is $\overline{u}$. If the agent rejects the contract the principal receives a payoff of zero.

Assume that the principal will want to offer the agent a contract such that the agent will accept (expected payoff is greater than $\overline{u}$). The principal's problem then is:

$$\max_{e \in \{e_L, e_H\},\ w(\pi)} \int_{\underline{\pi}}^{\overline{\pi}} (\pi - w(\pi)) f(\pi|e) \, d\pi$$

$$\text{s.t.} \quad \int_{\underline{\pi}}^{\overline{\pi}} v(w(\pi)) f(\pi|e) \, d\pi - g(e) \geq \overline{u}$$

Given that effort is observable, what is the best choice of $w(\pi)$ for each choice of e? Once that is known, what is the best choice of e?

If we split the first term into two pieces, $\int_{\underline{\pi}}^{\overline{\pi}} (\pi - w(\pi)) f(\pi|e) \, d\pi = \int_{\underline{\pi}}^{\overline{\pi}} \pi f(\pi|e) \, d\pi - \int_{\underline{\pi}}^{\overline{\pi}} w(\pi) f(\pi|e) \, d\pi$, we can see that choosing $w(\pi)$ to maximize $\int_{\underline{\pi}}^{\overline{\pi}} (\pi - w(\pi)) f(\pi|e) \, d\pi$ is the same as choosing $w(\pi)$ to minimize $\int_{\underline{\pi}}^{\overline{\pi}} w(\pi) f(\pi|e) \, d\pi$. Thus, we have:

$$\min_{w(\pi)} \int_{\underline{\pi}}^{\overline{\pi}} w(\pi) f(\pi|e) \, d\pi$$

$$\text{s.t.} \quad \int_{\underline{\pi}}^{\overline{\pi}} v(w(\pi)) f(\pi|e) \, d\pi - g(e) \geq \overline{u}$$

We can argue that the constraint always binds because why would the principal give the agent more than is needed to accept the contract? There may be reasons why that occurs in other situations, but in this specific model there is no reason for the principal to give the agent more than is required. The problem is a constrained optimization problem. Let γ be the multiplier on the constraint, so the agent's wage $w(\pi)$ for each $\pi \in [\underline{\pi}, \overline{\pi}]$ must satisfy:

$$\begin{aligned} -f(\pi|e) + \gamma v'(w(\pi)) f(\pi|e) &= 0 \\ \frac{1}{v'(w(\pi))} &= \gamma \end{aligned}$$

What happens if the agent is strictly risk averse? Then the risk-neutral principal just offers a fixed payment to the risk-averse agent based on the observable effort, and the principal fully insures the agent. Thus, the principal offers some fixed wage w_e^* such that $v(w_e^*) - g(e) = \overline{u}$. The wage w_e^* will depend on the effort level that is provided, with $w_{e_H}^* > w_{e_L}^*$ because it is more costly for the agent to exert high effort.[39]

What is the optimal choice of e? Recall that effort is observable, so the principal can specify e. The effort level is the one that maximizes expected profit minus wages:

$$\int_{\underline{\pi}}^{\overline{\pi}} \pi f(\pi|e) \, d\pi - v^{-1}(\overline{u} + g(e)).$$

The specific choice of e_H or e_L depends on both $f(\pi|e)$ and $g(e)$.

Proposition 7 *In the principal-agent model with observable effort, an optimal contract specifies that the agent choose the effort e^* that maximizes $\left[\int_{\underline{\pi}}^{\overline{\pi}} \pi f(\pi|e) \, d\pi - v^{-1}(\overline{u} + g(e))\right]$ and pays the agent a fixed wage $w^* = v^{-1}(\overline{u} + g(e))$. This contract is the uniquely optimal contract if $v''(w) < 0$ at all w.*

[39] If the agent is risk-neutral then any compensation scheme, including the fixed payment, is optimal as long as the expected wage payment is equal to $\overline{u} + g(e)$.

3.1.2 Unobservable effort

The optimal contract when effort is observable specifies an efficient effort choice and fully insures the agent against risk. When effort is not observable, these goals are in conflict because the principal pays a wage based on realized profit, not effort. Thus, the agent could exert high effort, get a bad profit draw, and be paid below his effort cost. When the agent is risk-neutral the principal can still achieve the same expected payoff as when effort is observable.

Risk-neutral agent Let $v(w) = w$. The optimal effort level e^* when effort is observable solves:

$$\max_{e \in \{e_L, e_H\}} \int_{\underline{\pi}}^{\overline{\pi}} \pi f(\pi|e)\, d\pi - g(e) - \overline{u}.$$

Now consider the case when effort is not observable. First, note that the principal can never do better when effort is unobservable (and unable to be specified by the principal) than when effort is observable (and able to be specified by the principal). If the principal could do better with unobservable effort, the principal could just specify the unobservable effort contract when effort is observable and allow the agent to choose effort.

Now let the principal specify $w(\pi) = \pi - \alpha$, where α is a fixed payment. Thus, the principal has essentially sold the project to the agent for α. Suppose the agent accepts. The agent then chooses e to maximize (remember the agent is now the owner of the project, so there is no issue with aligning incentives):

$$\int_{\underline{\pi}}^{\overline{\pi}} w(\pi) f(\pi|e)\, d\pi - g(e) = \int_{\underline{\pi}}^{\overline{\pi}} \pi f(\pi|e)\, d\pi - \alpha - g(e).$$

Note that e^* maximizes both $\int_{\underline{\pi}}^{\overline{\pi}} \pi f(\pi|e)\, d\pi - \alpha - g(e)$ (the unobservable payoff) and $\int_{\underline{\pi}}^{\overline{\pi}} \pi f(\pi|e)\, d\pi - g(e) - \overline{u}$ (the observable payoff) because the only difference between the two is the constants α and $\overline{u}$. So we have shown that the effort level is the same in the two problems.

Now, the agent will accept the contract $w(\pi) = \pi - \alpha$ as long as the agent's utility is at least $\overline{u}$, so:

$$\int_{\underline{\pi}}^{\overline{\pi}} \pi f(\pi|e^*)\, d\pi - \alpha - g(e^*) \geq \overline{u}.$$

Let α^* be the value of the fixed payment such that $\int_{\underline{\pi}}^{\overline{\pi}} \pi f(\pi|e^*)\, d\pi - \alpha^* - g(e^*) = \overline{u}$. Now, $\alpha^* = \int_{\underline{\pi}}^{\overline{\pi}} \pi f(\pi|e^*)\, d\pi - g(e^*) - \overline{u}$, which is the exact same expected payoff the principal had when effort was observable. There are no risk sharing problems when the agent is risk-neutral.

Risk-averse agent Now consider the case of a risk-averse agent. With risk-neutrality and unobservable effort we did not really utilize an incentive compatibility constraint because when the principal sold the project to the agent there was no need to align incentives. Formally:

$$\min_{w(\pi)} \int_{\underline{\pi}}^{\overline{\pi}} w(\pi) f(\pi|e)\, d\pi$$

$$\text{s.t.} \quad \int_{\underline{\pi}}^{\overline{\pi}} v(w(\pi)) f(\pi|e)\, d\pi - g(e) \geq \overline{u}$$

$$e \text{ solves } \max_{\widetilde{e}} \int_{\underline{\pi}}^{\overline{\pi}} v(w(\pi)) f(\pi|\widetilde{e})\, d\pi - g(\widetilde{e})$$

So the principal must now specify the contract in order to elicit the "correct" effort amount from the agent.

What wage function should the principal specify if he wants the agent to exert low effort e_L? He just offers $w_e^* = v^{-1}(\overline{u} + g(e_L))$. In order to induce the low effort level the principal simply specifies the fixed payment *as if effort is observable*. In this case, while effort is not observable, the principal knows that with this contract the agent will not choose high effort (unless the agent makes a mistake), so effectively the principal knows the effort level choice of the agent. This wage contract yields the same payoff to the principal in the unobservable effort case as when effort is observable, and we know that the principal can never earn more when effort is unobservable, so this contract must be a solution.

The contract to induce the agent to exert high effort is more interesting. As in the discrete payoff case above, the incentive compatibility constraint can be written as a comparison between exerting high effort and low effort, where the payoff to the agent exerting high effort must be greater than or equal to the payoff to the agent exerting low effort:

$$\int_{\underline{\pi}}^{\overline{\pi}} v\left(w\left(\pi\right)\right) f\left(\pi|e_H\right) d\pi - g\left(e_H\right) \geq \int_{\underline{\pi}}^{\overline{\pi}} v\left(w\left(\pi\right)\right) f\left(\pi|e_L\right) d\pi - g\left(e_L\right)$$

Now we have a constrained optimization problem with two inequality constraints. Let γ be the multiplier on the participation constraint and μ be the multiplier on the incentive compatibility constraint and we have:[40]

$$\begin{aligned}
\mathcal{L}\left(w, e, \gamma, \mu\right) &= \begin{array}{c} \int_{\underline{\pi}}^{\overline{\pi}} w\left(\pi\right) f\left(\pi|e_H\right) d\pi + \gamma\left(\overline{u} - \int_{\underline{\pi}}^{\overline{\pi}} v\left(w\left(\pi\right)\right) f\left(\pi|e_H\right) d\pi\right) \\ +\mu\left(\int_{\underline{\pi}}^{\overline{\pi}} v\left(w\left(\pi\right)\right) f\left(\pi|e_L\right) d\pi - g\left(e_L\right) - \int_{\underline{\pi}}^{\overline{\pi}} v\left(w\left(\pi\right)\right) f\left(\pi|e_L\right) d\pi + g\left(e_H\right)\right) \end{array} \\
0 &= \frac{\partial \mathcal{L}}{\partial w} \\
0 &= f\left(\pi|e_H\right) - \gamma v'\left(w\left(\pi\right)\right) f\left(\pi|e_H\right) + \mu\left(f\left(\pi|e_L\right) - f\left(\pi|e_H\right)\right) v'\left(w\left(\pi\right)\right) \\
f\left(\pi|e_H\right) &= \gamma v'\left(w\left(\pi\right)\right) f\left(\pi|e_H\right) - \mu\left(f\left(\pi|e_L\right) - f\left(\pi|e_H\right)\right) v'\left(w\left(\pi\right)\right) \\
f\left(\pi|e_H\right) &= v'\left(w\left(\pi\right)\right)\left[\gamma f\left(\pi|e_H\right) - \mu\left(f\left(\pi|e_L\right) - f\left(\pi|e_H\right)\right)\right] \\
\frac{1}{v'\left(w\left(\pi\right)\right)} &= \frac{\gamma f\left(\pi|e_H\right) - \mu\left(f\left(\pi|e_L\right) - f\left(\pi|e_H\right)\right)}{f\left(\pi|e_H\right)} \\
\frac{1}{v'\left(w\left(\pi\right)\right)} &= \frac{\gamma f\left(\pi|e_H\right) + \mu\left(f\left(\pi|e_H\right) - f\left(\pi|e_L\right)\right)}{f\left(\pi|e_H\right)} \\
\frac{1}{v'\left(w\left(\pi\right)\right)} &= \gamma + \frac{\mu\left(f\left(\pi|e_H\right) - f\left(\pi|e_L\right)\right)}{f\left(\pi|e_H\right)} \\
\frac{1}{v'\left(w\left(\pi\right)\right)} &= \gamma + \mu\left[1 - \frac{f\left(\pi|e_L\right)}{f\left(\pi|e_H\right)}\right]
\end{aligned}$$

Note that this result looks fairly similar to the observable effort case − in fact, it is the same, except for the $\mu\left[1 - \frac{f(\pi|e_L)}{f(\pi|e_H)}\right]$ term, which arises because of the presence of the incentive compatibility constraint. We know both γ and μ are positive, so both constraints bind. How? Consider $\mu = 0$. If that is true, then the payment is fixed at γ. For a fixed payment the agent will choose e_L because the disutility of e_H is greater than the disutility of e_L. To show $\gamma > 0$, suppose that the constraint is not binding, so $\gamma = 0$. Because $F\left(\pi|e_H\right)$ first-order stochastically dominates $F\left(\pi|e_L\right)$, there must exist an open set of profit levels $\Pi \subset [\underline{\pi}, \overline{\pi}]$ such that $\left[\frac{f(\pi|e_L)}{f(\pi|e_H)}\right] > 1$. If that is true, then $v'\left(w\left(\pi\right)\right) < 0$, which is a contradiction.

That γ and μ are positive leads to some interesting results. Suppose that $\widehat{w}$ is the fixed wage payment that leads to $\frac{1}{v'(\widehat{w})} = \gamma$. Then:

$$\begin{aligned}
w\left(\pi\right) &> \widehat{w} \text{ if } \frac{f\left(\pi|e_L\right)}{f\left(\pi|e_H\right)} < 1 \\
w\left(\pi\right) &< \widehat{w} \text{ if } \frac{f\left(\pi|e_L\right)}{f\left(\pi|e_H\right)} > 1
\end{aligned}$$

Intuitively, this result suggests that the principal is paying the agent more for outcomes that are statistically more likely to happen under e_H than under e_L. In class there was some discussion about whether or not $f\left(\pi|e_L\right) > f\left(\pi|e_H\right)$ for any profit levels. The thought was that high effort should lead to an increase in the likelihood of a particular profit level. While that statement is true for some profit levels, it cannot be true for all profit levels. Regardless of whether high effort or low effort is used, profits is always in the range

[40] Clearly there is a set of first-order conditions. As in MWG, we focus on one first-order condition.

of $[\underline{\pi}, \overline{\pi}]$. Thus, if $\overline{\pi}$ is more likely under high effort, some profit level $\widehat{\pi} < \overline{\pi}$ must be more likely under low effort.

What is interesting is that there are cases in which higher profits do not necessarily lead to higher wages (the wage structure is not monotonically increasing). If that is true, then $\frac{f(\pi|e_L)}{f(\pi|e_H)}$ must be decreasing in π. That means that as π increases, the likelihood of obtaining profit level π if effort is e_H relative to the likelihood if effort level is e_L must increase. While that seems logical, *first-order stochastic dominance does not guarantee it.* The text has an example of a distribution that first-order stochastically dominates another distribution, yet the optimal wage scheme is nowhere close to monotonic (see page 486). One final note is that if the principal wants to induce high effort in the unobservable case then he will have to pay the agent a higher wage than in the observable case. The higher wage compensates the agent for the risk that must be borne.

Note the effect that the unobservable case has on welfare. First, consider the case when the principal wants to induce effort e_L by the agent. We have seen that the optimal wage scheme under both is the same, so there is no welfare loss. Now, consider the principal inducing high effort. If high effort is optimal when observable, then either (1) the principal must compensate the agent more in order to induce the agent to exert high effort when effort is unobservable or (2) the principal may not find it optimal to induce high effort when effort is unobservable, thus leading the principal to induce low effort. Both lead to a welfare loss for the principal.

3.2 Revenue Equivalence Example

The math behind the concept of revenue equivalence. We are assuming that the environment is the SIPV-RN environment.

Let V_1 be the highest value and V_2 be the second highest value. Then the expected revenue of the 1^{st}-price auction is:

$$Revenue\left(1^{st} - price\right) = \frac{N-1}{N} E\left[V_1\right]$$

The expected revenue of the 2^{nd}-price auction is:

$$Revenue\left(2^{nd} - price\right) = E\left[V_2\right]$$

We will assume that there are the same number of bidders in each auction. We now need to know what $E\left[V_1\right]$ and $E\left[V_2\right]$ are in order to answer which of the auctions will generate more revenue. To do this we use the concept of an order statistic – basically, an order statistic tells us what the expected value of the k^{th} highest draw from a distribution will be given that we make N draws from the distribution. In our case, we are using the uniform distribution over the range 0 to 1. We find that the k^{th} highest value will be equal to:

$$\frac{N-k+1}{N+1}$$

Think about what this means. When there are 2 bidders, on average the highest value draw will be $\frac{2}{3}$, and on average the 2^{nd} highest value draw will be $\frac{1}{3}$. When there are 3 bidders, on average the highest value draw will be $\frac{3}{4}$, the 2^{nd} highest value draw will be $\frac{2}{4}$, and the 3^{rd} highest value draw will be $\frac{1}{4}$. Using this formula we have:

$$E\left[V_1\right] = \frac{N}{N+1}$$

$$E\left[V_2\right] = \frac{N-1}{N+1}$$

Now if we substitute these variables into our revenue for the 1^{st} and 2^{nd}-price auctions we get:

$$\begin{aligned} Revenue\left(1^{st} - price\right) &= \frac{N-1}{N} * \left(\frac{N}{N+1}\right) = \frac{N-1}{N+1} \\ Revenue\left(2^{nd} - price\right) &= \frac{N-1}{N+1} \end{aligned}$$

This means that the expected revenue from the 1^{st}-price auction is equal to $\frac{N-1}{N+1}$ and the expected revenue from the 2^{nd}-price auction is also equal to $\frac{N-1}{N+1}$. Thus, both auction formats are expected to generate the same revenue.

References

[1] Abdulkadiroğlu, Atila, Yeon-Koo Che, and Yosuke Yasuda (2011). Resolving Conflicting Preferences in School Choice: The "Boston Mechanism" Reconsidered. *American Economic Review* 101:1, 399-410.

[2] Abdulkadiroğlu, Atila, Parag A. Parthak, Alvin E. Roth, and Tayfun Sönmez (2005). The Boston Public School Match. *American Economic Review Papers and Proceedings* 95:2, 368-371.

[3] Abdulkadiroğlu, Atila and Tayfun Sönmez (2003). School Choice: A Mechanism Design Approach. *American Economic Review* 93:3, 729-747.

[4] Banks, Jeffrey, Mark Olson, David Porter, Stephen Rassenti, and Vernon Smith (2003). Theory, Experiment, and the Federal Communications Commission Spectrum Auctions. *Journal of Economic Behavior & Organization* 51, 303-350.

[5] Benoit, Jean-Pierre (1992). Scoring Reversals: A Major League Dilemma. *Social Choice and Welfare* 9, 89-97.

[6] Burden, Barry C. (2005). Minor Parties and Strategic Voting in Recent U.S. Presidential Elections. *Electoral Studies* 24, 603-618.

[7] Calsamiglia, Caterina, Guillaume Haeringer, and Flip Klijn (2010). Constrained School Choice: An Experimental Study. *American Economic Review* 100:4, 1860-1874.

[8] Chen and Sönmez (2006). School Choice: An Experimental Study. *Journal of Economic Theory* 127, 202-231.

[9] Cramton, Peter (1997). The FCC Spectrum Auctions: An Early Assessment. *Journal of Economics & Management Strategy* 6:3, 431-495.

[10] Cramton, Peter and Jesse A. Schwartz (2000). Collusive Bidding: Lessons from the FCC Spectrum Auctions. *Journal of Regulatory Economics* 17:3, 229-252.

[11] Dur, Umut, Robert G. Hammond, and Onur Kesten (2021). Sequential School Choice: Theory and Evidence from the Field and Lab. *Journal of Economic Theory*, forthcoming.

[12] Erdil, Aytek and Haluk Ergin (2008). What's the Matter with Tie-Breaking? Improving Efficiency in School Choice. *American Economic Review* 98:3, 669-689.

[13] Ergin, Haluk and Sönmez (2006). Games of School Choice Under the Boston Mechanism. *Journal of Public Economics* 90, 215-237.

[14] Fudenberg, Drew and Jean Tirole (1991). *Game Theory*. The MIT Press, Cambridge, MA USA.

[15] Gibbons, Robert (1992). *Game Theory for Applied Economists*. Princeton University Press, Princeton, NJ USA.

[16] Haeringer, Guillaume and Flip Klijn (2009). Constrained School Choice. *Journal of Economic Theory* 144, 1921-1947.

[17] Jones, Brian K. (2016). *The Regional Greenhouse Gas Initiative: A Three Article Study on the Decision to Join and Impacts of a Cap and Trade Market*. PhD dissertation, University of North Carolina at Charlotte.

[18] Kesten, Onur (2010). School Choice with Consent. *Quarterly Journal of Economics* 125:3, 1297-1348.

[19] Lien, Jaimie, Jie Zheng, and Xiaohan Zhong (2016). Preference Submission Timing in School Choice Matching: Testing Fairness and Efficiency in the Laboratory. *Experimental Economics* 19, 116-150.

[20] Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green (1995). *Microeconomic Theory.* Oxford University Press, New York, NY USA.

[21] Myerson, Roger (1979). Incentive Compatibility and the Bargaining Problem. *Econometrica* 47, 61-73.

[22] Pais, Joana and Ágnes Pintér (2008). School Choice and Information: An Experimental Study on Matching Mechanisms. *Games and Economic Behavior* 64, 303-328.

[23] Plott, Charles R. and Timothy C. Salmon (2004). The Simultaneous Ascending Auction: Dynamics of Price Adjustment in Experiments and in the UK3G Spectrum Auction. *Journal of Economic Behavior & Organization* 53, 353-383.

[24] Roth, Alvin E. (1982). The Economics of Matching: Stability and Incentives. *Mathematics of Operations Research* 7:4, 617-628.

[25] Roth, Alvin E. (1985). The College Admissions Problem Is Not Equivalent to the Marriage Problem. *Journal of Economic Theory* 36:2, 277-288.

[26] Roth, Alvin E. (2003). The Origins, History, and Design of the Resident Match. *Journal of the American Medical Association* 289:7, 909-912.

[27] Roth, Alvin E. and Elliott Peranson (1999). The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design. *American Economic Review* 89:4, 748-780.

[28] Weber, Roberto J. (1997). Making More from Less: Strategic Demand Reduction in the FCC Spectrum Auctions. *Journal of Economics & Management Strategy* 6:3, 529-548.

[29] Wolfstetter, Elmar (1999). *Topics in Microeconomics: Industrial Organization, Auctions, and Incentives.* Cambridge University Press, Cambridge, UK.